

---

# HEP Computing in the **Future**

Richard P. Mount

SLAC

*June 6, 2013*

## Summary: Follow the Money

---

- How Do We Spend Our Money?
- Classical Economics or What?  
(aka Distorting Factors)
- Unit Cost History of HEP Computing
- Murky Crystal Ball Gazing on Technology
- Murkier Crystal Ball Gazing on HEP Computing


# How do we spend our money?

---

- Hard numbers:
  - » US-ATLAS M&O spending on software and computing
- Educated guesses
  - » US spending on networking serving ATLAS
  - » US university and lab support for ATLAS computing
  - » US scientist effort on ATLAS software and computing operations
- Wild extrapolation
  - » Scale up to all US HEP and,
  - » Worldwide HEP


# US-Atlas Software and Computing M&O

## \$18.5M in 2015




# US-Atlas Software and Computing Total Cost Estimate


## \$27.0M in 2015




# Worldwide HEP Software and Computing Total Cost Guess \$200M in 2015



# Hardware Costs (BNL T1 average 2013-2017)



# Hardware-Related Costs (BNL T1 average 2013-2017)





# Classical Economics or What?

---

- Strong focus on CPU as the implied cost driver and figure of merit
- Political/funding requirement for a distributed system
  - » A centralized system might divide unit costs by a factor 2
  - » But it might divide funding by a factor 4
- Some real benefits of the distributed system
  - » Involves scientists worldwide and distributes expertise
  - » Eases access to major opportunistic resources
  - » Maybe it doesn't really cost us all that much

# Classical Economics or What?


---


- Effort is “free” (up to a rigid limit) in some national systems
- Much of the network is “free”
- Space/Power/Cooling (even sysadmins) are “free” at many centers
- Tape is on its way out (just as it was for the last 15 years), so don’t make it part of any serious planning (This is the prevailing wisdom)

# Unit Cost History of HEP Computing

---

- Simple filter – stuff that I have bought (with some network numbers from Harvey Newman)
- Added Titan and Mira





# Murky Crystal Ball – Rotating Disks

---

- Market is shrinking
  - » Consumer market is shrinking rapidly
  - » Enterprise market is growing
  - » Cheap, low performance, reliable (tape-killer) disks are disappearing
- Disk technology has a long way to go if likely market volumes justify the investments
- Slower evolution of TB/\$ - doubling every 3 years?

# Murky Crystal Ball – Tape

---

- Has been at death's door for more than a decade
- Everybody hates it, (strange technology, people-intensive, maintenance ransom, ...) but
- Everybody needs archival storage
- The threat from disk is receding
- Make tape an integral part of the planning

# Murky Crystal Ball – Solid State Storage

---

- High market demand (and limited fabrication capacity)
  - » Current prices are high relative to “equilibrium” market
  - » Costs about 10 times rotating disk
  - » May go down to 2 to 3 times over the next 5 to 10 years
  - » Not clear how strongly the market will drive the introduction of “better” technologies than flash



# Murky Crystal Ball – CPU

---

- We hit the GHz wall five years ago
  - » Didn't really hurt us much
- We may not quite hit the feature size wall (too few atoms in a feature) in the next decade, but it is coming
- I, personally, have been in painful collision with the power consumption wall several times in the last decade
- The “doubling every 1.3 years” growth has stopped. But we will continue to get more transistors per \$.
- The transistors will be slower, but will do more work per Watt
- CPU units/\$ may grow a little faster than disk (IF we can use the transistors).

# Murky Crystal Ball – WAN

---

- Not quite an open market – regulations, cartels.
- e.g. No transatlantic cables being pulled
- Guess that bits/s per \$ will continue to grow on average to keep pace with average storage+CPU costs

# Future of HEP Computing

---

- Distributed Computing will persist
  - » Funding politics will not change
  - » Our distributed approach eases the use of opportunistic or pay-as-you-go resources that don't belong to us. (for example HPC)
  - » Need to focus on lowering the people cost
- Costs will continue to be dominated by people
- Need to focus on overall cost optimization  
(Some money sources may appear to be CPU obsessed – don't let that distract us)

# Future of HEP Computing

---

- Engineering software for future architectures
  - » No clear picture of how the market will drive architectures, but
  - » Event-level parallelism is unlikely to be enough
  - » Need more:
    - HEP software effort (significant cost)
    - Collaboration with CS and other sciences
    - Coordination
  - » Looks like an ongoing need more than a project.

# Future of HEP Computing

---

- Storage
  - » Multi layer (solid state/disk/tape and successors)
  - » Cost optimized combination
  - » Optimization will depend critically on application-sensitive cache-management
  - » We may have to provide much of the software ourselves

# Future of HEP Computing

---

- Crystal ball is dark on
  - » How long will C++/Python dominate?
    - But we must use software technologies that make our grad students employable.
  - » How do we maintain high-quality software effort in the face of new physics data?
  - » Are systems of the complexity of ATLAS Distributed Computing sustainable long-term?

# Summary

---

“Prediction is very difficult, especially about the future”

*(Attributed to nearly everybody from Bohr to Berra)*